

Maximum-entropy moment-closure for stochastic systems on networks

Tim Rogers

School of Physics & Astronomy, The University of Manchester, M13 9PL, UK

E-mail: tim.rogers@manchester.ac.uk

Abstract. Moment-closure methods are popular tools to simplify the mathematical analysis of stochastic models defined on networks, in which high dimensional joint distributions are approximated (often by some heuristic argument) as functions of lower dimensional distributions. Whilst undoubtedly useful, several such methods suffer from issues of non-uniqueness and inconsistency. These problems are solved by an approach based on the maximisation of entropy, which is motivated, derived and implemented in this article. A series of numerical experiments are also presented, detailing the application of the method to the Susceptible-Infective-Recovered model of epidemics, as well as cautionary examples showing the sensitivity of moment-closure techniques in general.

Submitted to: *JStat*

1. Introduction

A great many problems of interest in statistical physics and applied mathematics may be modelled as a system of stochastic variables whose interactions form a network. The earliest such models were inspired by questions in chemistry and typically imposed a lattice interaction structure (see [1] and references therein for many examples), though as the applications of these models have diversified to include everything from finance [2] to ecology [3], the range of network topologies under consideration has expanded also. Whilst of great benefit to the real-world relevance of networked models, the complexity of the network structures now in vogue causes serious complications to any attempted mathematical analysis.

In the absence of a network structure, the behaviour of well-mixed stochastic models may often be very closely described by systems of differential equations, which have proven to be indispensable tools in their analysis. It is a long-standing problem to reproduce this success in the field of networked models. The primary obstacle to this endeavor is the existence of correlations between spatially separated parts of the networked model. In mathematical terms, one finds that differential equations describing the joint distribution of some number of variables will typically depend on

joint distributions of higher dimension, leading to a never ending expansion in the number of equations required ‡.

For as long as this problem has been recognised, there have been approximation schemes suggested to deal with it. Often trading under the name ‘moment-closure’ these techniques amount to choosing a method by which to guess at higher dimensional distributions, based on knowledge of the lower dimensional distributions for which one wishes to write differential equations.

The simplest non-trivial example of a moment-closure scheme is the pair approximation, proposed for ecological applications in [6, 7, 8], but already with a long history in statistical physics (see, for example [9, 10, 11]). In this approximation, the state of the whole system is reduced to the joint probability of pairs of variables which neighbour each other in the network. The joint distribution of three connected variables is then approximated by one of two particular functions, depending on whether they form a triangle or a line in the network. This approximation has been widely applied and is successful for a reasonably broad range of models and parameters [3, 12, 13, 14], particularly in those of adaptive networks [15, 16, 17, 18]. There is, however, a mathematical problem with the pair approximation as it is usually applied to triangles. The formula given in [8, 19] for the joint distribution of three variables forming a triangle is in general not a probability distribution at all (it is not normalised) and is typically inconsistent with the known marginal distribution of pairs. We will address this issue in detail in Sections 2.4 and 3.2.

Moving beyond the pair approximation, more complex schemes have been proposed based on joint probabilities of three or more variables; examples include those of references [14, 20, 21]. The moment-closure methods employed in most such works are reminiscent of Bayes nets [22], being based on the assumption of some form of conditional independence, allowing the factorisation of high dimensional joint probabilities into smaller objects. Therein lies a problem of non-uniqueness: depending on the local configuration of the network, there may be several such factorisations to choose from and no logical motivation for making one choice over another. For a discussion of this problem in relation to an approximation based on triples, see [14].

This article will develop a moment-closure approach based on the maximisation of entropy. Put simply, when the need arises to approximate a high dimensional distribution based on knowledge of some lower dimensional marginals, it is wise to avoid introducing any additional, unmotivated, bias; this is achieved by choosing the distribution of maximum entropy subject to the known constraints. The approach has the distinct advantages of providing a unique and consistent approximation scheme for distributions of any dimension, and numerical experiments demonstrate considerable success in certain models. There is, however, the aesthetic downside that often the resulting approximation cannot be written as an explicit function of the lower dimensional distributions. In these circumstances, one must deal with it either implicitly

‡ Alternatively, consideration of the time evolution of macroscopic observables in disordered systems on networks leads to a subtly different closure problem which has been addressed in, e.g. [4, 5].

or numerically; some appropriate methods are outlined in Section 2.3. It should be noted that the idea of applying maximum-entropy to obtain moment-closure approximations is not unique to this article, having been discussed recently for continuous space models in physics [23] and ecology [24]. The emphasis here is on the application of the method to networked models, and the relationship to existing alternative schemes.

To connect with other approaches, the maximum-entropy formalism is applied to the pair approximation in Section 2.4. In dealing with the case of three sites in a row the pair and maximum-entropy approximations agree, whilst for triangles we are able to suggest an alternative to the inconsistent choice usually made. The performance of this proposed alternative approximation is checked numerically in Section 3.2, with the surprising result that, whilst it is considerably better at predicting the joint distribution on triangles, the output from numerical integration of the moment-closed differential equations is notably worse. Further numerical investigation reveals that this phenomenon can be attributed to a bizarrely fortunate cancellation of error which occurs within the usual pair approximation when applied to networks with many triangles.

To bring to a close the numerical investigations of the article, an experiment is proposed to test the robustness of moment-closure methods in general. The basic philosophy of moment-closure rests on the assumption that having a good approximation to the required higher dimensional joint distributions will endow the moment-closed system of differential equations with high predictive power. Despite its central importance to the method, the nature of the relationship between accuracy of approximation and predictive success has not been well studied. In Section 3.3 the results of a simple experiment are presented in which higher dimensional distributions are not approximated but rather measured from a simulation. It is shown that the application of a small systematic error in one direction or another induces a rather large discrepancy in the output of numerical integration of the moment-closed differential equations. This result suggests that the quality of predictions made by moment-closure approximations may be considerably more fluid than one might hope.

2. Maximum-entropy approximation

2.1. Description of the problem

We consider a discrete time stochastic process \S defined on a simple graph $\mathcal{G} = (V, E)$. For notational simplicity, we label the sites by natural numbers: $V = \{1, \dots, N\}$, where N is the size of the graph. If $(i, j) \in E$ we say that i and j are neighbours in network, and we write ∂i for the set of all neighbours of i . Upon each site i a time dependent variable X_i is placed, taking values in a finite set of states. In each timestep, a site i is chosen at random and the state of the variable X_i is updated. The new value is taken to depend only on the old value and those of its neighbours.

\S No generality is lost here by not discussing continuous time processes as both formulations have the same first order behaviour under system size expansion.

A great number of processes of interest in statistical physics may be modelled according to this broad description. To fix ideas, we will refer throughout this article to the paradigmatic example of the Susceptible-Infected-Recovered (SIR) model of epidemics. Some number N of individuals are arranged in a network, with the presence of an edge indicating the possibility of infectious contact. The possible states of an individual are denoted S , I and R , for susceptible, infected and recovered, respectively. Per timestep a site i is chosen at random and the variable X_i updated according to: if $X_i = I$ then we set $X_i = R$ with probability γ ; if $X_i = R$ then we set $X_i = S$ with probability α ; if $X_i = S$ then a neighbour $j \in \partial i$ is chosen at random and if $X_j = I$ then we set $X_i = I$ with probability β . The last reaction of this list models the transmission of the disease, which is restricted to pairs of neighbouring sites and thus the spread of the disease may be strongly dependent on the geometry of the network.

Rather than always working with the detailed information of the exact state of each variable in a stochastic process, it is usual to consider macroscopic observables – statistical measures which summarise some aspect of the state of the system at a given time. The simplest example is the distribution of states of a randomly chosen site, we write

$$P(x) = \frac{1}{|V|} \sum_{i \in V} \mathbb{I}\{X_i = x\},$$

where \mathbb{I} is the indicator function. Objects of this type are very useful in giving an overview of the behaviour of the system, for example, the progress of an SIR epidemic may be monitored by plotting the time evolution of $P(I)$, that is, the fraction of the population currently infected.

To gain more detailed information about the underlying process, the above construction may be generalised to arbitrary arrangements of variables in the network. Specifically, for a given (assumed small) graph g , let $P_g(\mathbf{x})$ denote the fraction of induced subgraphs $h \subset \mathcal{G}$ which are isomorphic to g and whose site variables take exactly the states $\mathbf{x} = (x_1, x_2, \dots)$. Formally, we write

$$P_g(\mathbf{x}) = \frac{\sum_{h \subset \mathcal{G}} \mathbb{I}\{h \cong g\} \prod_{i \in h} \mathbb{I}\{X_i = x_i\}}{\sum_{h \subset \mathcal{G}} \mathbb{I}\{h \cong g\}}, \quad (1)$$

where \cong denotes graph isomorphism. A simple example is given by the joint distribution of pairs of neighbouring variables, setting g to be the graph of two connected sites, we may define

$$P(x, y) := P_g(x, y) = \frac{1}{|E|} \sum_{(i, j) \in E} \mathbb{I}\{X_i = x, X_j = y\}.$$

We will refer to distributions of the type (1), which are measured directly from an underlying stochastic process, as empirical. Note that empirical distributions will change with time, though we have not made this explicit in the notation.

The usual methodology of analyses based on objects of the type $P_g(\mathbf{x})$ is to move from regarding them as measurements of a stochastic system to treating them as dynamical variables in their own right. This is achieved by rescaling time by a factor of N^{-1} and taking the limit $N \rightarrow \infty$ (the limit of large networks), in which it is assumed that $P_g(\mathbf{x})$ will approach a smooth non-random function, satisfying a differential equation which can be derived from the details of the underlying stochastic process. The complicating factor in this approach when applied to networked stochastic systems is that generally the rate of change of $P_g(\mathbf{x})$ will depend not only on those variables which contribute to $P_g(\mathbf{x})$, but also on their neighbours, leading to a system which is not closed. There is more than one way to perform this analysis and different sets of differential equations may be found (for example, compare the approaches to SIR epidemics taken in references [14] and [25]), however, a relatively general situation is described by a collection of differential equations of the form

$$\frac{d}{dt}P_g(\mathbf{x}_g) = \sum_{\mathbf{y}} P_G(\mathbf{y})\Gamma(\mathbf{x}_g, \mathbf{y}), \quad (2)$$

where G is the graph containing those sites in g and their neighbours, and the functions $\Gamma(\mathbf{x}_g, \mathbf{y})$ describe probability flows between states, derived from the microscopic dynamics. In this form, such equations are of very little use as they say that the rate of change of P_g depends on the higher dimensional distribution P_G , which is not part of the system. Moment-closure techniques seek to address this problem by suggesting possible ways to express P_G approximately as a function of P_g . In the next subsection, we will derive the maximum-entropy approach to moment-closure.

2.2. Derivation

Consider a graph G , covered by a collection \mathbb{G} of subgraphs. To each site $i \in G$ a random variable X_i is attached: write \mathbf{X} for the vector containing these variables, and \mathbf{X}_g for the projection of that vector onto the set of sites also included in the subgraph $g \in \mathbb{G}$. The X_i are assumed to have a joint distribution function P_G which is *unknown*. Suppose that we have knowledge of the marginal distribution on each of the subgraphs, given by

$$P_g(\mathbf{x}_g) := \mathbb{P}[\mathbf{X}_g = \mathbf{x}_g] = \sum_{\mathbf{x}_i : i \notin g} P_G(\mathbf{x}).$$

In the context of moment-closure, the subgraph distributions P_g are the level at which we wish to close our system of differential equations and G is the larger graph containing all variables which influence the time evolution of a particular P_g . For a simple example, the pair approximation corresponds to the situation in which G is a single site plus its neighbours, and \mathbb{G} is the collection of edges in G .

The core of the problem is as follows: to close our system of differential equations at the level of the P_g , it is necessary to approximate the unknown higher-dimensional distribution P_G , guided only by knowledge of the marginal distributions P_g .

It is correct to demand that the function \tilde{P}_G approximating P_G must of course be a probability distribution itself, and that its marginals \tilde{P}_g should agree with the known distributions P_g , that is:

$$\text{for all } \mathbf{x}, \text{ and for every } g \in \mathbb{G}, \text{ we have } \sum_{x_i : i \notin g} \tilde{P}_G(\mathbf{x}) = P_g(\mathbf{x}_g). \quad (3)$$

Aside from these requirements, in a general setting there is no *a priori* reason to assume that \tilde{P}_G should take any particular form. It can therefore be argued [26] that the most sensible choice for \tilde{P}_G is the distribution which has maximum entropy subject to the above constraint. Put simply, to make any other choice for \tilde{P}_G would introduce an additional bias which would require further justification.

Accepting this reasoning, the calculation is straightforward. It will suffice to satisfy (3), as the requirement that \tilde{P}_G is a bona fide distribution will then hold automatically. The entropy of a distribution P is given by

$$E(P) = - \sum_{\mathbf{x}} P(\mathbf{x}) \log P(\mathbf{x}).$$

To perform the constrained maximisation of this quantity, we introduce one Lagrange multiplier $\lambda_g(\mathbf{x}_g)$ per subgraph $g \in \mathbb{G}$ and projected state vector \mathbf{x}_g . We thus seek extrema of the functional

$$\Lambda[P] = - \sum_{\mathbf{x}} P(\mathbf{x}) \log P(\mathbf{x}) + \sum_g \sum_{\mathbf{x}_g} \lambda_g(\mathbf{x}_g) \left(P_g(\mathbf{x}_g) - \sum_{x_i : i \notin g} P(\mathbf{x}) \right).$$

Differentiating, we find

$$\frac{\partial \Lambda[P]}{\partial P(\mathbf{x})} = -1 - \log P(\mathbf{x}) - \sum_g \lambda_g(\mathbf{x}_g),$$

and thus the maximum entropy approximation for P_G is given by

$$\tilde{P}_G(\mathbf{x}) = \exp \left\{ -1 - \sum_g \lambda_g(\mathbf{x}_g) \right\}.$$

We see that \tilde{P}_G factorises over subgraphs; there exist functions Q_g such that

$$\tilde{P}_G(\mathbf{x}) = \prod_g Q_g(\mathbf{x}_g), \quad (4)$$

where the equations for Lagrange multipliers imply that the functions Q_g collectively satisfy

$$\text{for all } g \text{ and } \mathbf{x}_g, \quad P_g(\mathbf{x}_g) = \sum_{\mathbf{y}} \delta(\mathbf{x}_g - \mathbf{y}_g) \prod_h Q_h(\mathbf{y}_h). \quad (5)$$

Note that the normalisation of the functions $P_g(\mathbf{x}_g)$ was not accounted for during the implementation of the Lagrange multipliers, meaning that there are several redundant degrees of freedom in the solution set to (5) which should be treated with care. Indeed, it is straightforward to see that (4) is invariant under certain multiplicative transforms of the Q_g , however, this is only a problem if one intends to work directly with the Q_g in a numerical context.

2.3. Practical implementation

Although it is known [26] that (provided the supplied marginals P_g are consistent) the maximum entropy distribution \tilde{P}_G always exists and is unique, only in certain circumstances can it be written neatly as a function of the marginals. We can characterise these situations by considering the graph \mathcal{H} , with sites given by the elements of \mathbb{G} (that is, the subgraphs of G for which marginals are supplied), and edge set given by the rule that $g \sim h$ if and only if g and h share at least one site. If \mathcal{H} is a tree, then

$$\tilde{P}_G(\mathbf{x}) = \prod_g \frac{P_g(\mathbf{x}_g)}{\sqrt{\prod_h P_{g \cap h}(\mathbf{x}_{g \cap h})}}, \quad (6)$$

where $P_{g \cap h}$ is the marginal distribution of variables contained in both g and h , and we set $P_\emptyset \equiv 1$ for consistency.

For a general setting in which \mathcal{H} is not a tree, there is no algebraic expression for \tilde{P}_G in terms of the P_g . Numerical solution of the problem is straightforward however, for example by a simple process of iterative replacement we describe now. Starting from a uniform distribution $P^{(0)}$, one repeatedly cycles through the subgraphs $g \in \mathbb{G}$, at each step setting

$$P^{(n+1)}(\mathbf{x}) = P_g(\mathbf{x}_g) \frac{P^{(n)}(\mathbf{x})}{\sum_{x_i: i \notin g} P^{(n)}(\mathbf{x})}.$$

As $n \rightarrow \infty$, the distributions $P^{(n)}$ converge to \tilde{P}_G . Proof of the convergence of this procedure may be found in [27], where it is employed in the entirely different context of quantifying distributional complexity by analysis of interaction structure.

The final consideration in the practical application of maximum-entropy moment-closure is the inclusion of the approximation in the differential equations it is intended to close. As discussed in the introduction, for a given microscopic dynamics there often exists more than one set of ODEs purporting to approximate the time evolution of the system. We consider the case that one wishes to close a system of differential equations of the form (2) which we assume are derived from some underlying stochastic process. In the fortunate situation that the graph \mathcal{H} of subgraph overlaps is a tree, one may simply apply equation (6) to close the system by setting

$$\begin{aligned} \frac{d}{dt} P_g(\mathbf{x}_g) &= \sum_{\mathbf{y}} \Gamma(\mathbf{x}_g, \mathbf{y}) P_G(\mathbf{y}) \\ &\approx \sum_{\mathbf{y}} \Gamma(\mathbf{x}_g, \mathbf{y}) \prod_{g'} \frac{P_{g'}(\mathbf{y}_{g'})}{\sqrt{\prod_h P_{g' \cap h}(\mathbf{y}_{g' \cap h})}}. \end{aligned} \quad (7)$$

Away from this special case one only has access to \tilde{P}_G via numerical methods, and no such simple expression is possible. However, if, as is very often the case, one intends to treat the differential equations numerically, then the lack of a simple closure formula is not much of a penalty. At each step of the numerical solution of the differential equations, the numerical value of \tilde{P}_G may be recovered efficiently via the iterative replacement scheme described above.

Alternatively, the equations can always be closed at the level of the functions Q_g , with the resulting system given by, for all g and \mathbf{x}_g ,

$$\sum_{\mathbf{y}} \left(\prod_h Q_h(\mathbf{y}_h) \right) \left[\Gamma(\mathbf{x}_g, \mathbf{y}) - \delta(\mathbf{x}_g - \mathbf{y}_g) \sum_h \frac{d}{dt} \log Q_h(\mathbf{y}_h) \right] = 0. \quad (8)$$

As a result of the under-specification of the Lagrange multipliers discussed earlier, there are some extraneous degrees of freedom in this system of differential equations, though these can be accounted for either by holding constant a suitable number of values $Q_g(\mathbf{x}_g)$, or by imposing additional normalisation constraints on the Q_g to prevent divergence.

To summarise, in the general case for which equation (7) does not apply, one has the choice of working with \tilde{P}_G numerically, or directly with the Q_g via (8).

2.4. Example - pair mean field

Consider the SIR model on a regular graph of degree k . In [25], the progression of an SIR epidemic was modelled by differential equations equivalent to

$$\begin{aligned} \frac{d}{dt} P(S, S) &= -2 \frac{\beta}{k} P(S, S, I) \\ \frac{d}{dt} P(S, I) &= \frac{\beta}{k} \left(P(S, S, I) - P(I, S, I) - P(S, I, I) \right) \\ \frac{d}{dt} P(S, R) &= -\frac{\beta}{k} P(R, S, I) + \gamma P(S, I) \\ \frac{d}{dt} P(I, I) &= 2 \frac{\beta}{k} \left(P(I, S, I) + P(S, I, I) \right) - 2 \gamma P(I, I) \\ \frac{d}{dt} P(I, R) &= \frac{\beta}{k} P(R, S, I) + \gamma \left(P(I, I) - P(I, R) \right). \end{aligned} \quad (9)$$

In these expressions k denotes the average degree (number of neighbours) of the sites in the graph, $P(x, y)$ is probability distribution of a pair of adjacent variables, and

$$P(x, y, z) = (1 - \phi) P_R(x, y, z) + \phi P_T(x, y, z), \quad (10)$$

where $P_R(x, y, z)$ is the distribution of three sites in a row, $P_T(x, y, z)$ is the distribution of three sites forming a triangle, and ϕ is a measure of the frequency of triangles in the underlying graph. To close these equations at the level of $P(x, y)$, we are required to find expressions for $P_R(x, y, z)$ and $P_T(x, y, z)$. Note that although these equations involve distributions of dimensions two and three and thus are not of the same form as (2), maximum-entropy moment-closure may still be applied.

Considering $P_R(x, y, z)$ first, we are interested in the situation illustrated in Figure 1 – we wish to compute the maximum-entropy approximation \tilde{P}_R , under the constraint that each of the pairs defined by subgraphs a and b has the known marginal distribution $P(x, y)$. From equation (4), we see that there exist functions Q_a and Q_b such that

$$\tilde{P}_R(x, y, z) = Q_a(x, y) Q_b(y, z).$$

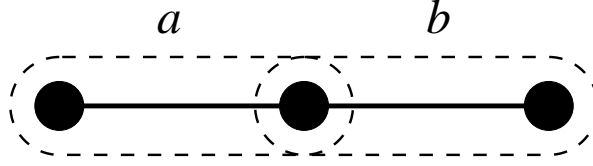


Figure 1: The graph R consisting of three sites arranged in a row, covered by subgraphs a and b given by linked pairs of sites.

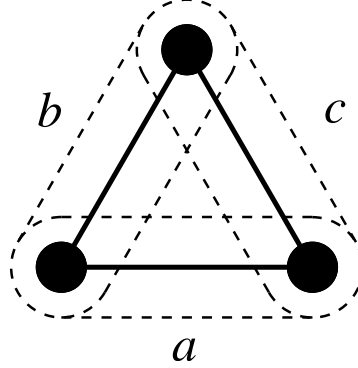


Figure 2: The graph T consisting of three sites arranged in a triangle, covered by subgraphs a , b and c , given by linked pairs of sites.

The constraints (5) then give first for subgraph a ,

$$P(x, y) = P_a(x, y) = \sum_z Q_a(x, y) Q_b(y, z) = Q_a(x, y) \sum_z Q_b(y, z)$$

$$\Rightarrow Q_a(x, y) = \frac{P(x, y)}{\sum_z Q_b(y, z)},$$

and then for subgraph b ,

$$P(y, z) = P_b(y, z) = \sum_x Q_a(x, y) Q_b(y, z) = Q_b(y, z) \sum_x \frac{P(x, y)}{\sum_z Q_b(y, z)}$$

$$\Rightarrow Q_b(y, z) = \frac{P(y, z)}{P(y)} \sum_z Q_b(y, z).$$

Combining the two yields

$$\tilde{P}_R(x, y, z) = \frac{P(x, y) P(y, z)}{P(y)}. \quad (11)$$

Alternatively, we might observe that there are only two subgraphs and we may then apply the simple form (6) directly to obtain the same result.

Equation (11) is precisely the usual pair approximation to the joint distribution of three variables connected in a row, so we see that the maximum-entropy approach agrees with the usual methods in this case. Turning attention now to the case of the triangle (illustrated in Figure 2), we find that the situation is not so straightforward.

The graph of subgraph overlaps in this case is not a tree and thus there is no simple algebraic form for \tilde{P}_T . In [19, 8] an approximation is suggested equivalent to

$$P_T(x, y, z) \approx F(x, y, z) := \frac{P(x, y)P(y, z)P(z, x)}{P(x)P(y)P(z)}. \quad (12)$$

This expression is sometimes [19, 28] attributed to Kirkwood [29]. Unfortunately, for almost all P_T this function is not a distribution itself and does not agree with the known marginals on pairs. In defense of this expression, it is usually clearly advertised as an approximation made for convenience, and besides we will see later that it holds up remarkably well when used to numerically integrate equations (9), despite being inaccurate in general. Note that (12) is precisely the result one would obtain from naively (incorrectly) applying (6) in this situation.

For numerical applications, the maximum-entropy approximation to P_T may easily be computed using the iterative replacement algorithm outlined earlier. Such an approach may however be unappealing to those who would like a neat closed system of ODEs, whose terms do not depend on quantities which must be computed numerically. In that case, one might choose to halt the iterative replacement process after some number of steps. For example, three steps will yield the approximate form

$$P_T(x, y, z) \approx \hat{P}_T(x, y, z) = \frac{P(x, y)P(y, z)P(z, x)}{\sum_w P(x, w)P(z, w) \frac{P(x)}{P(w)}}, \quad (13)$$

which is a viable alternative to (12) with the advantages firstly of being a genuine probability distribution and secondly of agreeing with the marginal distribution on the pair (x, y) . As will be reported in the next section, numerical experiments comparing the approximations (12) and (13) have surprising results.

3. Numerical experiments

3.1. An epidemic on the square lattice

As a first test of the utility of the maximum-entropy approach, we attempt to predict the behaviour of the SIR model on a square lattice by numerically integrating a moment-closed system of equations. The pair approximation on a square lattice has already been considered in [30], so we will develop a higher order theory based on clusters of five connected sites.

Let g be the graph of a single site and its neighbours in the square lattice – see Figure 3a for an illustration. To close the equations for SIR epidemics at the level of P_g will require approximation of the distribution P_G on the graph G of all sites within distance two of a central site (Figure 3b). This bigger graph contains five copies of g (Figure 3c), which are used as the basis of the maximum-entropy approximation.

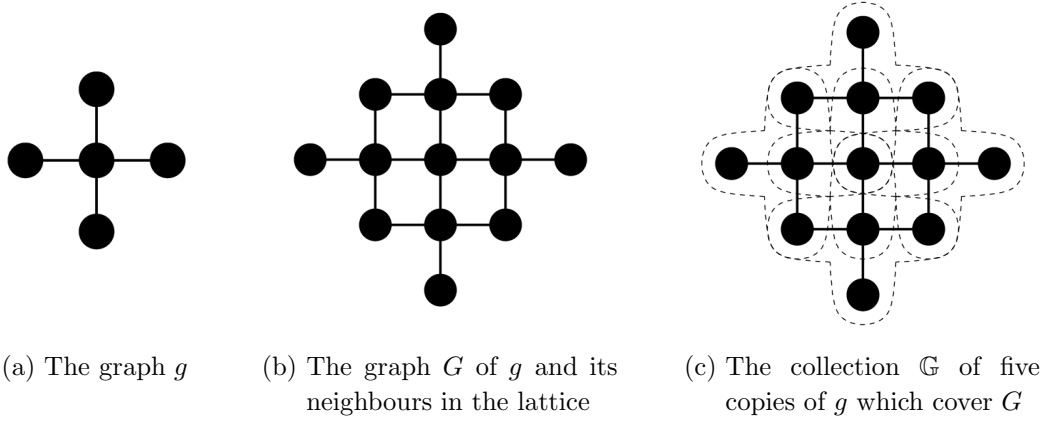


Figure 3: Developing a theory for the SIR model on a square lattice which is closed at the level of P_g (a site and its neighbours) requires estimating the distribution of variables in P_G , which itself contains five copies of g .

As per the earlier discussion, one formulation for the moment-closed system of differential equations is given by

$$\frac{d}{dt}P_g(\mathbf{x}_g) = \sum_{\mathbf{y}} \tilde{P}_G(\mathbf{y}) \Gamma(\mathbf{x}_g, \mathbf{y}), \quad (14)$$

however, to compute the coefficients $\Gamma(\mathbf{x}_g, \mathbf{y})$ one must enumerate the possible changes to each $h \in \mathbb{G}$ which could occur given the configuration \mathbf{y} of the sites in G . Explicitly,

$$\Gamma(\mathbf{x}_g, \mathbf{y}) = \sum_{h \in \mathbb{G}} \left[\mathbb{P}(\mathbf{y}_h \mapsto \mathbf{x}_g \mid \mathbf{y}) - \mathbb{I}\{\mathbf{x}_g = \mathbf{y}_h\} \left(1 - \mathbb{P}(\mathbf{x}_h \mapsto \mathbf{x}_h \mid \mathbf{y}) \right) \right],$$

where \mapsto is used as a shorthand for a change taking place in one timestep.

Note that for SIR dynamics there are three possible states per site which, given that there are five sites in g and thirteen in G , means that the system (14) contains 243 equations and moreover each one has 1594323 terms in its sum. The solving of this system is most certainly a job for the computer - see [31] for a discussion of this problem in relation to the pair approximation.

Instead of attempting to list the equations themselves, a better understanding can be gained through an examination of the steps involved in numerically integrating them. Starting from an initial condition in which each state of each site is chosen independently at random,

$$P_g(\mathbf{x}_g) = \prod_{i \in g} P(x_i),$$

we apply the Euler forward method of numerical integration. At each step, time is moved on by a small amount $t \mapsto t + \Delta t$, and the following computation is performed:

- (i) Initially, for each \mathbf{x}_g , set

$$\frac{d}{dt}P_g(\mathbf{x}_g) = 0$$

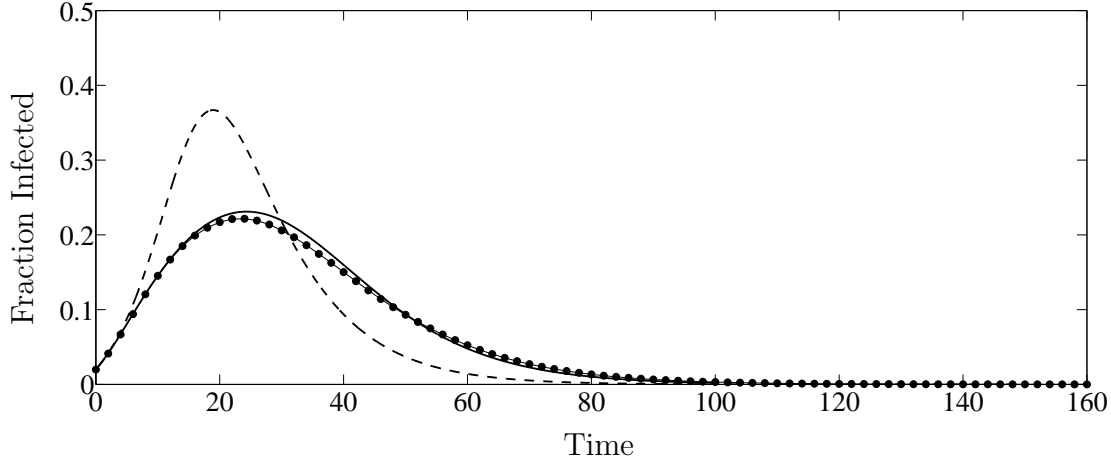


Figure 4: Comparison of simulation results (circles) of SIR dynamics on a square lattice with predictions of the pair approximation (dashed line) and the five site maximum-entropy approximation (solid line) described in the text. The disease parameters are $\beta = 0.6$, $\gamma = 0.1$ and $\alpha = 0$.

(ii) For every configuration \mathbf{y} of variables in G

- (a) Compute the maximum-entropy approximation $\tilde{P}_G(\mathbf{y})$ by iterative replacement
- (b) Compute the probability δ of a change to the state of the central variable in this configuration, write \mathbf{y}' for the resulting configuration
- (c) For each $h \in G$ put

$$\frac{d}{dt}P_g(\mathbf{y}_h) \mapsto \frac{d}{dt}P_g(\mathbf{y}_h) - \delta\tilde{P}_G(\mathbf{y})$$

and

$$\frac{d}{dt}P_g(\mathbf{y}'_h) \mapsto \frac{d}{dt}P_g(\mathbf{y}'_h) + \delta\tilde{P}_G(\mathbf{y})$$

(iii) For each \mathbf{x}_g , update

$$P_g(\mathbf{x}_g) \mapsto P_g(\mathbf{x}_g) + \Delta t \left(\frac{d}{dt}P_g(\mathbf{x}_g) \right).$$

The results of this process for two different parameter choices are presented in Figures 4 and 5. In each case a single random simulation of a discrete time SIR epidemic on a toroidal square lattice of size $N = 10^6$ was performed according to the microscopic dynamics outlined in Section 2.1.

For Figure 4 the parameters for infection, recovery and return to susceptibility were chosen to be $\beta = 0.6$, $\gamma = 0.1$ and $\alpha = 0$, respectively. Alongside the result from the simulation are those obtained by the numerical integration scheme for the cluster approximation described above, as well as the pair approximation of [30] for comparison. As is clearly visible in the figure, the cluster approximation provides a considerable improvement to the accuracy of the prediction.

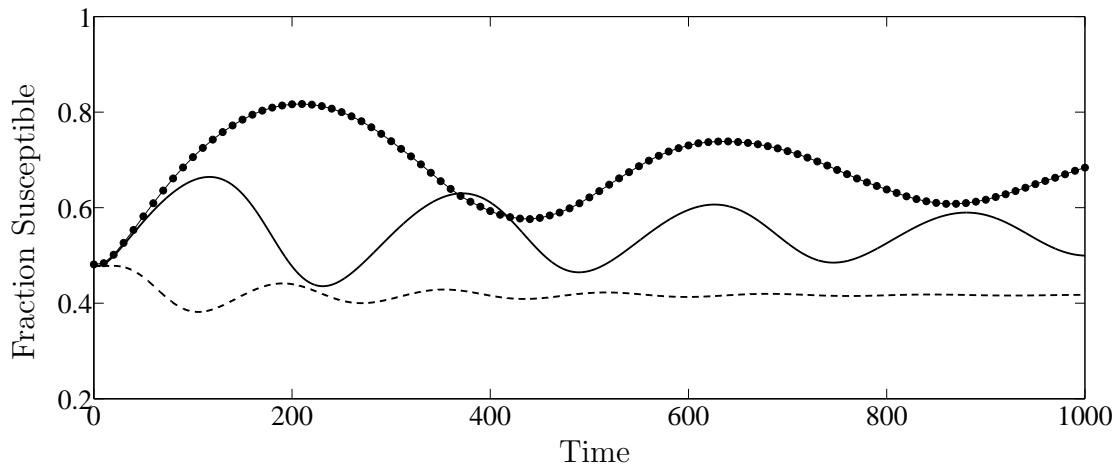


Figure 5: Comparison of simulation results (circles) of SIR dynamics on a square lattice with predictions of the pair approximation (dashed line) and the five site maximum-entropy approximation (solid line) described in the text. The disease parameters are $\beta = 0.6$, $\gamma = 0.1$ and $\alpha = 0.01$.

An identical set-up was used for Figure 5, with the exception that return to susceptibility was allowed with probability $\alpha = 0.01$. This change has the effect of moving the system into an oscillatory regime in which both the pair approximation and the cluster approximation are dramatically less accurate. The failure of moment-closure methods in this regime can be attributed to the development of long range correlations in the model which appear to play a central role in controlling the frequency of oscillations observed. It is worth noting, though, that for many practical applications the goal of moment-closure methods is to quantify the steady states of the system (and their stability), rather than to accurately reproduce transient features such as the decaying oscillations observed here. In this regard, the cluster approximation is again a clear improvement over the pair approximation.

3.2. Comparison of pair approximations

In the previous section, we saw that the maximum-entropy approximation for three sites in a row agreed with the usual pair approximation for such a configuration. For a triangular arrangement of sites, however, the situation is rather more complicated. The commonly cited approximation F , given in equation (12) does not satisfy the requirement of being a probability distribution. An alternative expression \hat{P}_T was suggested here in equation (13), found as the partial implementation of maximum-entropy, and may possibly offer an improvement.

The following pair of numerical experiments are designed to give a casual assessment of the performance of these approximations. Firstly, the outputs are compared directly to the empirical distribution they are intended to approximate, secondly, the expressions F and \hat{P}_T are employed in the numerical integration of the differential equations (9) to

see how success or failure in the first test may affect the predictive power of the system they close.

The approximations in question deal with triangles, so a network with plenty of these will provide an appropriate setting for the numerical tests. We choose a small-world network constructed as follows: N sites are arranged in a circle and each is connected to its four nearest neighbours, each edge then has one end randomly rewired with probability δ ($= 0.1$ for the present case). Random graphs of this type have frequently been used as models of social networks and thus the spread of disease on such networks is interesting in its own right. For the simulations used here a single graph of this type of size $N = 1000$ was generated in this way. The frequency of triangles in this graph was measured at $\phi = 0.355$. A sample of 1000 independent runs of SIR dynamics were then performed on this graph, using an initial condition of 20 infective sites chosen at random, and parameters for infection, recovery and return to susceptibility given by $\beta = 0.8$, $\gamma = 0.2$ and $\alpha = 0$, respectively.

For the first test, at each tenth of a timestep (this is possible since time is rescaled by a factor of N^{-1} in the microscopic dynamics) the empirical distribution $P_T(x, y, z)$ of three sites arranged in a triangle was computed using equation (1). The empirical marginal distribution $P(x, y)$ of pairs of edges *which are involved in a triangle* was computed by summing $P_T(x, y, z)$ over one argument.

From equations (9) and (10), one sees that the quantities in need of approximation are $P_T(S, S, I)$, $P_T(I, S, I)$ and $P_T(R, S, I)$, which between them describe the distribution of the state of the third site in a triangle in which the others form a susceptible-infected pair. The shaded areas in the top plot of Figure 6a shows the empirical conditional density $P_T(x | S, I) = P_T(x, S, I)/P(S, I)$ for $x = S, I, R$, obtained from the simulation. Since they form a probability distribution, they necessarily sum to one. For the central plot of Figure 6a, the approximate conditional density $\hat{P}_T(x | S, I)/P(S, I)$ at each data point was computed using equation (13), with the empirical pair distribution $P(x, y)$ taken as the input. Finally, the lower plot of the same figure shows the result F of equation (12), computed in the same way.

The difference in the accuracy of the approximations F and \hat{P}_T shown in Figure 6a is striking. Whilst \hat{P}_T appears to give a reasonably close approximation to the empirical conditional distribution $P_T(\cdot | S, I)$, the results of F are a dramatic underestimate, particularly of the fraction of susceptibles.

In the production of Figure 6a, the empirical pair distribution $P(x, y)$ was used for as the basis for F and \hat{P}_T . If instead the approximations are employed for their intended purpose of closing the system of equations (9), the pair distribution used will evolve over time according to these equations and, as a result, errors may quickly become compounded to produce an inaccurate prediction of the progress of the disease. In Figure 6b, the time evolution of the fraction of infected sites as averaged over the simulations (circles) is shown alongside the predictions resulting from system (9) closed using approximations F (dashed line) and \hat{P}_T (solid line).

As is clearly visible in the figure, the success of equation \hat{P}_T as an approximation

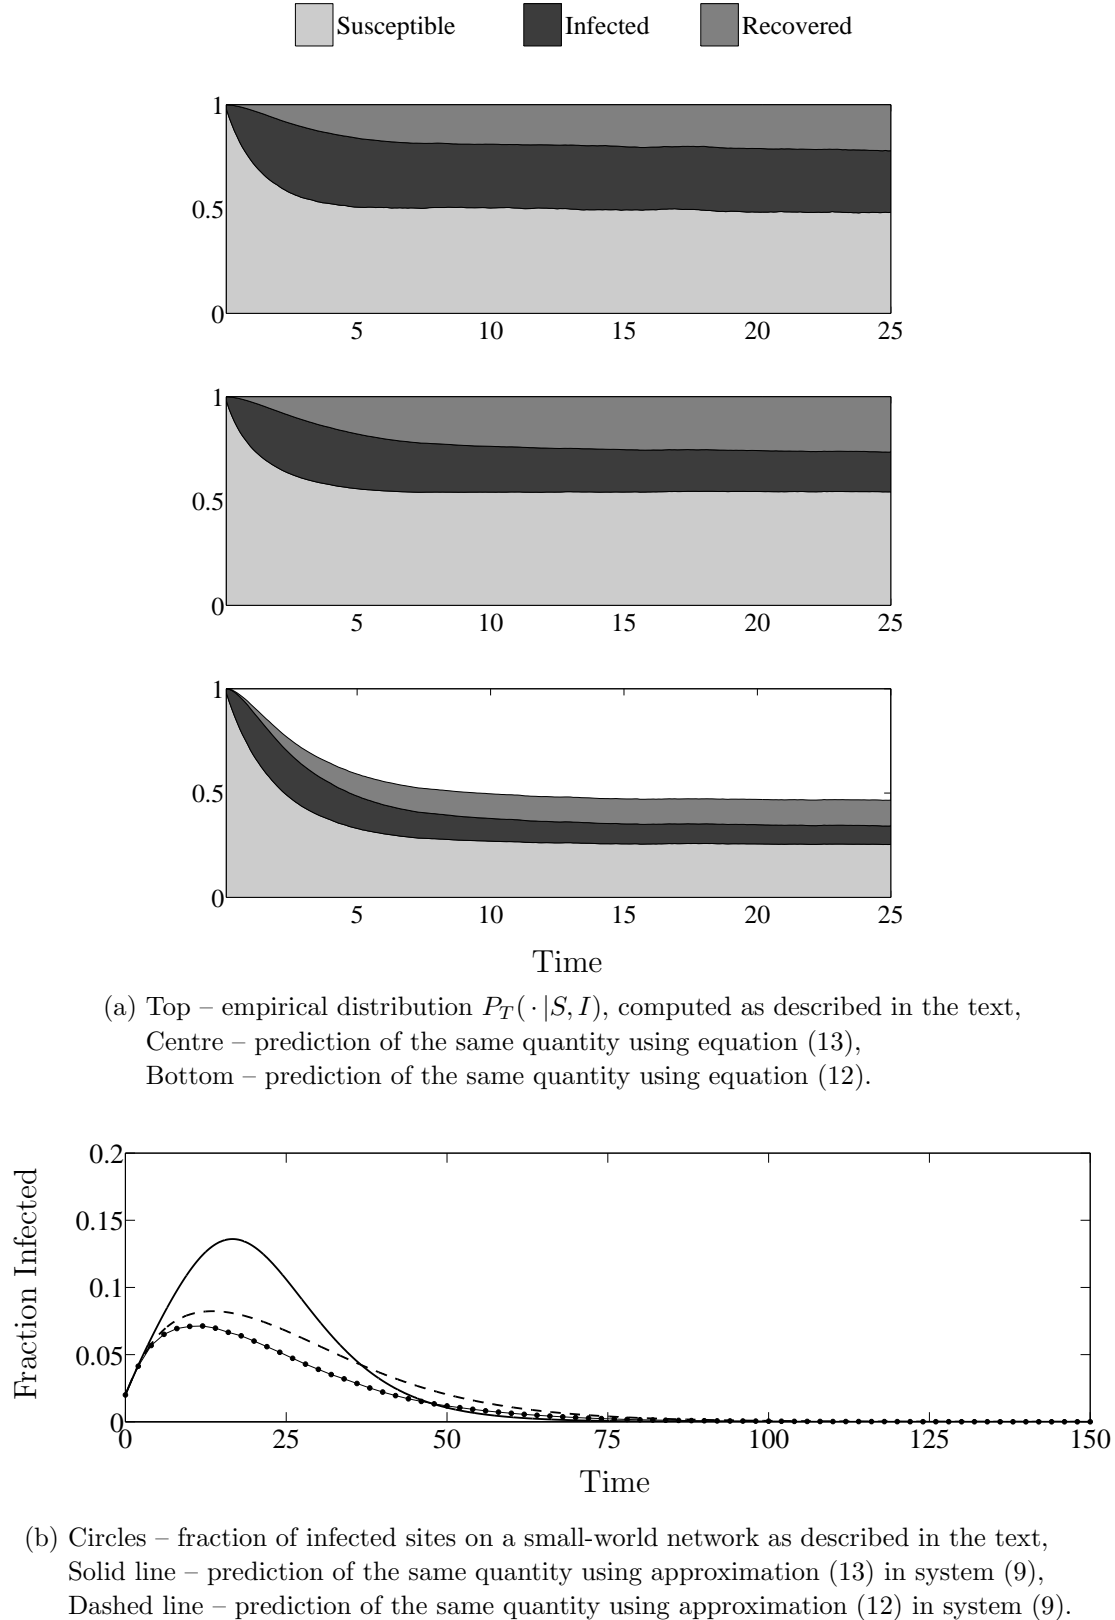


Figure 6

to the distribution P_T has not translated into a good prediction when used to close the system of differential equations. In contrast, the approximation F is relatively successful when used in this way, providing a much closer fit to the data from numerical simulations.

That the dramatic failure of F in reconstructing the density P_T does not result in a poor prediction of the progress of the disease is surprising and requires explanation. Note that both approximation schemes make use of the same estimate \tilde{P}_R for the distribution P_R of three sites in a row, which here accounts for around 65% of the estimate of $P(x, y, z)$, see equation (10). Performing a similar numerical experiment to assess the accuracy of this shared component reveals that (for this graph and parameter choice at least) it provides a substantial overestimate of the conditional density of susceptibles $P_R(S|S, I)$. In contrast, Figure 6a shows that approximation F gives a large underestimate of the quantity $P_T(S|S, I)$, which has the effect of lowering the overall estimate of $P(S, S, I)$, acting to cancel out most of the erroneous overestimation introduced by \tilde{P}_R . Since the approximation \hat{P}_T provides a much better estimation of $P_T(S|S, I)$, the error in the shared component \tilde{P}_R is carried over unchecked, resulting in a poorer overall performance.

To summarise, it appears that the success of the approximation scheme proposed in [19, 25] may in fact be largely due to a fortunate cancellation caused by two poor approximations giving large errors in opposite directions. As the experiments of this subsection have demonstrated, improving one part of the scheme actually has the counterproductive effect of destroying this curious symmetry of errors, resulting in a worse performance overall.

3.3. Sensitivity to errors in approximation

The central premise of moment-closure methods, including the maximum-entropy approach discussed in this article, is that having access to a ‘good’ approximation scheme to higher order correlations will translate into a high accuracy of prediction from a system of differential equations employing them. It is therefore sensible to ask: precisely how good is ‘good’? To put this query in concrete terms, the following numerical experiment tests the sensitivity of moment-closure to errors in the approximation.

Consider the following system of differential equations describing the time evolution of the fraction of susceptible, infected and recovered individuals in a standard SIR epidemic:

$$\begin{aligned}\frac{d}{dt}P(S) &= -\beta P(S)P(I|S) + \alpha P(R) \\ \frac{d}{dt}P(S) &= \beta P(S)P(I|S) - \gamma P(I) \\ \frac{d}{dt}P(S) &= \gamma P(I) - \alpha P(R).\end{aligned}\tag{15}$$

Here the term $P(I|S)$ denotes the conditional probability that a randomly chosen neighbour of a randomly chosen susceptible site is infected, so for example $\beta P(S)P(I|S)$

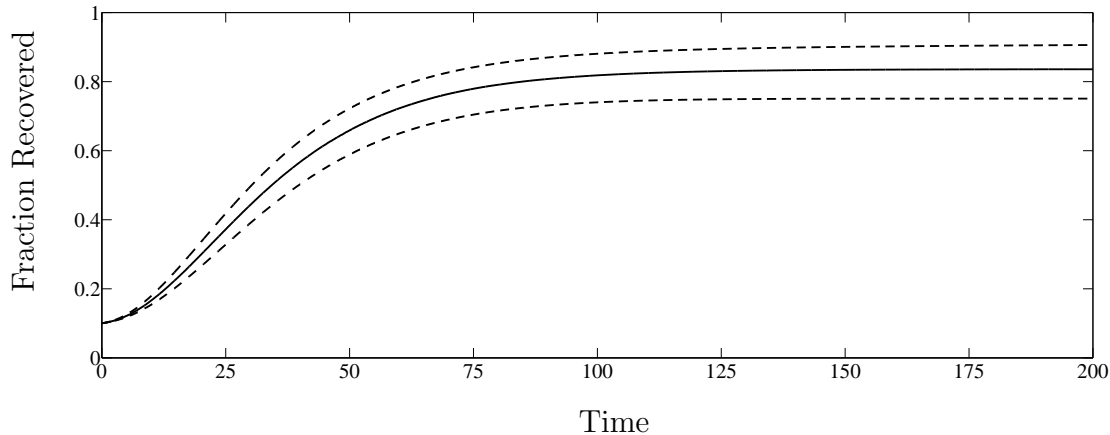


Figure 7: Solid line – fraction of recovered sites measured from a single run of SIR dynamics on a toroidal square lattice of size $N = 10^6$. Dashed lines – high and low results from numerical integration of differential equation system (15), with the pair distribution modified but within $\varepsilon = 0.001$ in KL-divergence from the empirical value.

gives the probability that a susceptible site is chosen for update and is then infected by a neighbour. These equations give an exact description (on average and in the limit $N \rightarrow \infty$) of the progress of an epidemic, however they are not closed as they depend on the unknown quantity $P(I|S)$.

Now consider the following numerical experiment. SIR dynamics are performed on a large square lattice and at each micro-timestep the empirical conditional probability $P(I|S)$ is measured directly from the simulation. Simultaneously with this process, the system of differential equations (15) is numerically integrated using the numerical value of $P(I|S)$ taken from the empirical measurement at the same moment in time. Since the equations (15) are exact and $P(I|S)$ is determined from the simulation, the output from the simulation and the numerical integration will be almost indistinguishable – this fact is easily verified on the computer.

By repeating this process, but now replacing $P(I|S)$ in the numerical integration by a slightly perturbed value, one will gain some insight into the robustness of the system (15) to errors in approximation of $P(I|S)$.

The results of such an experiment are shown in Figure 7. A single run of SIR dynamics was performed on a toroidal square lattice of size $N = 10^6$, with parameters $\beta = 0.6$, $\gamma = 0.1$ and $\alpha = 0$. The high and low perturbations of $P(I|S)$ were obtained by adjusting this value up or down as far as possible whilst keeping the Kullback Leibler divergence^{||} from the empirical distribution $P(\cdot|S)$ less than $\varepsilon = 0.001$.

For most practical purposes, a KL-divergence of $\varepsilon = 0.001$ is extremely low and would represent a very close approximation to the true distribution. As demonstrated in Figure 7, however, a such small range of error in the distribution can still result in

^{||} KL-divergence is a measure (not a metric) of the difference between two probability distributions, defined in the discrete case by $D(P_1||P_2) = \sum_x P_1(x) \log[P_1(x)/P_2(x)]$.

considerable uncertainty in the resulting prediction of system (15), as the error is quickly compounded over the course of the epidemic. Even with such a precise approximation to $P(I|S)$, the final size of the epidemic may be over- or under-estimated by around 10%.

On first glance it may seem unfair to choose extremal errors in $P(I|S)$, however, it should be noted that, as we saw in the previous subsection, errors in moment-closure approximations may often be systematic rather than random and hence such a choice is quite realistic.

4. Conclusion

Moment-closure methods are vital tools in the analysis of stochastic systems with network structure, allowing one to make mathematical predictions about the behaviour of models previously only accessible to simulation. The success of such methods is certainly not in question, however, it is mitigated by several issues:

- (i) the most common moment-closure method (the pair approximation) is inconsistent in its treatment of triangles
- (ii) higher order approximation schemes usually involve an arbitrary and unmotivated choice of some conditional independence structure
- (iii) moment-closed systems of differential equations (particularly those of higher order) are computationally intensive to solve
- (iv) the insertion of an approximation scheme into a system of differential equations obfuscates the effects of the accuracy of the approximation.

The work presented in this article provides an effective solution to the first two issues on this list. By following a systematic approach based on maximisation of entropy, one is able to obtain approximation schemes of any desired size and complexity, without the need for any arbitrary or inconsistent assumptions.

In those situations in which a conditional independence structure genuinely is present, such as the pair approximation of three sites in a row, it is recovered by maximum-entropy and simple algebraic equations for the approximation may be derived. In more general situations where no such simple equations exist, the maximum-entropy approximation may be computed efficiently by a process of iterative replacement. Resorting to numerical approaches is certainly an aesthetic failing but it is not a practical one as moment-closed systems of differential equations are already computationally intensive to solve and the use of the maximum-entropy approximation causes very little additional burden. Indeed, numerical computation of the closure terms is all that is required for analysis using software packages such as AUTO [32], and so the maximum-entropy method may be straightforwardly implemented there.

To investigate the application of maximum-entropy moment-closure, issue (ii) above was considered in some detail. The usual pair approximation to triangles suffers from the mathematical drawbacks of neither being properly normalised, nor agreeing with

the known marginal distribution on pairs. The maximum-entropy approximation by definition does not suffer either of these problems, however, it cannot be written down in a neat closed form. As a middle ground, an alternative pair approximation to triangles was obtained by halting the iterative replacement algorithm after three steps. In numerical experiments based on the SIR model on a small world network, this alternative scheme vastly outperformed the usual approximation in the job of predicting the joint distribution of three sites forming a triangle. However, when coupled with the approximation for three sites in a row and used to close a system of differential equations, the results of the new approximation did not agree with those of numerical simulations, whilst making the less accurate approximation to triangles somehow resulted in a better fit to the data. Further numerical investigation revealed that this unusual effect is due to a fortuitous cancellation of error between the pair approximations for rows and triangles, whose balance was upset by improving one half of the approximation scheme but not the other.

The phenomenon that an improved approximation scheme can lead to worse results when used to close a system of differential equations raises often overlooked questions about the reliability and sensitivity of moment-closure methods in general. As the simple numerical experiment presented in Section 3.3 shows, an approximation which is highly accurate but suffers from a very small systematic bias can yield results with errors in prediction which are several orders of magnitude larger. This suggests that predictions made by moment-closure methods (whose biases we do not know) should be viewed as being somewhat flexible.

Moving forward, there is considerable scope to apply the methods outlined in this article to the study of stochastic systems on networks. It would be particularly interesting to take a maximum-entropy approach to the study of moment closure methods in adaptive networks [15, 16, 17, 18] and stochastic corrections [14].

Acknowledgments

Thanks to Alan McKane, Tobias Galla and Joe Challenger for useful discussions and advice. Funding from the EPSRC under grant number EP/H02171X/1 is gratefully acknowledged.

- [1] Ziman J M 1979 *Models of disorder: the theoretical physics of homogeneously disordered systems* (Cambridge University Press)
- [2] May R M, Levin S A and Sugihara G 2008 *Nature* **451** 893–5
- [3] Dieckmann U, Law R and Metz J 2000 *The Geometry of Ecological Interactions* Cambridge Studies in Adaptive Dynamics (Cambridge University Press)
- [4] Mozeika A and Coolen A C C 2008 *Journal of Physics A: Mathematical and Theoretical* **41** 115003
- [5] Mozeika A and Coolen A C C 2009 *Journal of Physics A: Mathematical and Theoretical* **42** 195006
- [6] Matsuda H, Ogita N, Sasaki A and Sato K 1992 *Progress of Theoretical Physics* **88** 1035–1049
- [7] Satulovsky J and Tomé T 1994 *Physical Review E* **49** 5073–5079
- [8] Keeling M J, Rand D A and Morris A J 1997 *Proceedings of the Royal Society of London. Series B: Biological Sciences* **264** 1149–56
- [9] Kikuchi R 1951 *Physical Review* **81** 988–1003
- [10] Dickman R 1986 *Physical Review A* **34** 4246–4250
- [11] Tomé T and de Oliveira M 1989 *Physical Review A* **40** 6643–6646
- [12] Levin S A and Durrett R 1996 *Proceedings of the Royal Society of London. Series B: Biological Sciences* **351** 1615–21
- [13] de Aguiar M, Rauch E and Bar-Yam Y 2003 *Physical Review E* **67**
- [14] Rozhnova G and Nunes A 2010 *The European Physical Journal B* **74** 235–242
- [15] Gross T, DLima C and Blasius B 2006 *Physical Review Letters* **96**
- [16] Gross T and Kevrekidis I G 2008 *Europhysics Letters* **82** 38004
- [17] Zschaler G, Traulsen A and Gross T 2010 *New Journal of Physics* **12** 093015
- [18] Demirel G, Prizak R, Reddy P N and Gross T 2011 *The European Physical Journal B* 1–8
- [19] Matsuda H 2000 *Physical Review E* **62** 3096–3102
- [20] Szabó G, Szolnoki A and Izsák R 2004 *Journal of Physics A: Mathematical and General* **37** 2599–2609
- [21] Petermann T and De Los Rios P 2004 *Journal of theoretical biology* **229** 1–11
- [22] Pearl J 1988 *Probabilistic reasoning in intelligent systems: networks of plausible inference* (Morgan Kaufmann)
- [23] Singer A 2004 *The Journal of chemical physics* **121** 3657–66
- [24] Raghil M, Hill N A and Dieckmann U 2010 *Journal of mathematical biology* 1–49
- [25] Keeling M J 1999 *Proceedings of the Royal Society of London. Series B: Biological Sciences* **266** 859–67
- [26] Jaynes E 1957 *Physical Review* **106** 620–630
- [27] Csiszár I and Shields P C 2004 *Information theory and statistics: a tutorial* (Now Publishers Inc)
- [28] Filipe J 2003 *Mathematical Biosciences* **183** 15–35
- [29] Kirkwood J G 1942 *The Journal of Chemical Physics* **10** 394
- [30] de Souza D R and Tomé T 2010 *Physica A: Statistical Mechanics and its Applications* **389** 1142–1150
- [31] Alexander M E and Kobes R 2010 *Preprint: arXiv 1007.2883* 23
- [32] Doedel E J, Champneys A R, Fairgrieve T F, Kuznetsov Y A, Sandstede B and Wang X Auto 97: Continuation and bifurcation software for ordinary differential equations (with homcont)